

# DAOS v2.8 Features Demo

Ken Cain, Tom Nabarro, Wei Li

21 May 2026

# Agenda

- 
- 01** Feature: self\_heal properties (automatic recovery)

---

  - 02** Feature: Interactive rebuild control

---

  - 03** Recorded demo: Interactive rebuild control

---

self\_heal / Auto Recovery



# DAOS Self-Healing Problem Statement

- Component failure triggers rebuild to restore redundancy
- Automation not always desirable
- Want policies to gain control over recovery
- Coordination improves resource utility and stability



# Self-heal policy syntax

```
[daos]$ dmg system get-prop self_heal
```

Name	Value
----	-----
Self-heal policy for the system (self_heal)	exclude;pool_exclude;pool_rebuild

```
[daos]$ dmg system set-prop self_heal:"exclude;pool_exclude"  
system set-prop succeeded
```

To be used in conjunction with pool-level policies

```
[daos]$ dmg pool get-prop p1 self_heal
```

```
Pool p1 properties:  
Name          Value  
----          -  
Self-healing policy (self_heal) exclude;rebuild
```

```
[daos]$ dmg pool set-prop p1 self_heal:"none"  
pool set-prop succeeded
```



# Key Points

- Distinct rank memberships
- System-level policy in 2.8
- Needs no engines or pool services running
- Pool query lists policy settings
- Eval command triggers actions
- Evolution continues toward predictable automation



# System-Level self\_heal example

Show control over rebuild eliminating some unnecessary activity.

DAOS system with single pool, 8 ranks and default properties.

1. Set system property to **exclude;pool\_exclude** to prevent pool rebuild
2. Perform full system stop and subsequently start all but 2 ranks
  - Pool remains in idle rebuild but ranks are excluded from system and pools
3. Set property to **exclude;pool\_exclude;pool\_rebuild** to enable rebuild
  - Eval command triggers rebuild on pool



# Step 1

Set system property **exclude;pool\_exclude** to prevent rebuild

```
[daos]$ dmg system set-prop self_heal:"exclude;pool_exclude"  
system set-prop succeeded
```

```
[daos]$ dmg pool get-prop p1 self_heal
```

Pool p1 properties:

Name	Value
----	-----
Self-healing policy (self_heal)	exclude;rebuild

```
[daos]$ dmg pool query p1
```

```
Pool 68174afc-ce59-42b6-8580-14407cedde9d, ntarget=128, disabled=0, leader=6, version=1, state=Ready
```

```
...
```

```
rebuild disabled on pool due to [system] policy
```



## Step 2

Perform full system stop and start all but 2 ranks

- Pool remains in idle rebuild but ranks are excluded from system and pools

```
[daos]$ dmg system stop
```

```
Rank  Operation Result
```

```
----  -
```

```
[0-7] stop      OK
```

```
[daos]$ dmg system start -r0-5
```

```
Rank  Operation Result
```

```
----  -
```

```
[0-5] start      OK
```

```
[daos]$ dmg pool query p1
```

```
Pool 68174afc-ce59-42b6-8580-14407cedde9d, ntarget=128, disabled=32, leader=5, version=35, state=TargetsExcluded
```

```
Pool health info:
```

```
- Disabled ranks: 6-7
```

```
- Rebuild done, 0 objs, 0 recs
```

```
- Data redundancy: degraded
```

```
...
```



# Step 3

Choice:

- Restore data redundancy after failure on startup either with or without failed ranks
- EITHER repair/reintegrate OR eval command triggers rebuild on pool

```
[daos]$ dmg system set-prop self_heal:"exclude;pool_exclude;pool_rebuild"  
system set-prop succeeded
```

```
[daos]$ dmg system self-heal eval  
System self-heal eval request succeeded
```

```
[daos]$ dmg pool query p1  
Pool 68174afc-ce59-42b6-8580-14407cedde9d, ntarget=128, disabled=32, leader=5, version=35, state=TargetsExcluded  
Pool health info:  
- Disabled ranks: 6-7  
- Rebuild busy, 123 objs, 1234 recs  
- Data redundancy: degraded  
... >>>  
- Rebuild done, 0 objs, 0 recs  
- Data redundancy: normal  
...
```



# Interactive Rebuild Control



# Commands / Behavior, Rebuild State Query

- (New) Per-pool and system-wide control:

```
dmg pool <label> rebuild stop      dmg pool <label> rebuild start
dmg system rebuild stop            dmg system rebuild start
```

- `rebuild stop` is “one-shot” (stops only the *current* rebuild)
  - After stop, pool targets left in a rebuild-able state, for subsequent `rebuild start`
  - New rebuilds can be triggered by subsequent `dmg` command, fault reactions, pool service leader switch, ...

- (Existing) `dmg pool query Rebuild State` – new output for effects of a `rebuild stop`

```
Rebuild stopping (state=busy, status=-2027)      # -2027 == -DER_OP_CANCELED
                                                  # initial state shortly after rebuild stop command
Rebuild busy, <m> objs, <n> recs                 # briefly during Fail_reclaim (cleanup operation)
Rebuild stopped (state=idle, status=-2027)      # final state for a stopped rebuild
```



# Interactive Rebuild Motivations

- Context: Stop or Prevent (via `self_heal`) Rebuilds:
  - default `self_heal`: exclusion done, rebuilds running – can use `rebuild stop` if manual recovery preferred
  - `self_heal` without `pool_rebuild`: only exclusion done – can perform manual rebuilds
- Useful for Flexible Recovery Options – on exclusion, decide whether / when to rebuild
  - Do nothing for a while (allow pool to stay “degraded” / allow applications to run without competing with rebuild)
  - Recovery actions (possibly paced – performing in batches of pools vs. all at once)
    - `rebuild start` command(s) (if deciding to repair for the exclusions)
    - `reintegrate` command(s) (if deciding to directly reintegrate after quick resolution of original issue)
- Useful if a given rebuild fails and retries repeatedly
  - Manually stop, so network/CPU resources aren’t wasted ; start after fixing the underlying issue
- Useful if a long-running rebuild is suspected to be stuck (e.g., due to an unforeseen DAOS bug)



# Interactive Rebuild Control

(Recorded Demo)



# Demo – Prevent Rebuilds, Manually Perform Recovery

- System with Multiple Pools (2 for this demo – p1 and p2)
- Prevent Rebuilds – configure system `self_heal without pool_rebuild`
- Engines A,B lost – only exclusions, no rebuild (due to system `self_heal`) – see with `dmg pool query`
- Interactive Recovery
  - `dmg pool rebuild start p1` (wait until `dmg pool query` confirms rebuild is finished)
  - `dmg pool rebuild start p2` (wait until `dmg pool query` confirms rebuild is running)
  - Decide to change approach for pool p2 – choose direct reintegration:
    - `dmg pool rebuild stop p2` (wait until `dmg pool query` confirms this rebuild is stopped)
    - `dmg pool reintegrate -ranks=A,B p2` (wait until `dmg pool query` confirms this rebuild is finished)

**Thank You**



# References

- self\_heal documentation: [https://docs.daos.io/master/admin/self\\_healing](https://docs.daos.io/master/admin/self_healing)
- Interactive Rebuild Control documentation: [https://docs.daos.io/master/admin/rebuild\\_controls](https://docs.daos.io/master/admin/rebuild_controls)